Chapter 5

Coupling models of crustal deformation and mantle convection with a computational framework[1]

5.1 Abstract

Crust and mantle exhibit contrasting compositions and thus distinct rheologies, making it challenging to model their combined mechanical behaviors. A code coupling technique is explored in order to solve multi-material and multi-physics problems. We adapt Pyre, a framework allowing distinct codes to exchange variables through shared interfaces, to the coupling of a Lagrangian and an Eulerian code. Motivated by computationally challenging problems like crust-mantle interaction, we couple an explicit code for the solution of crustal deformation for the mechanics of lithosphere with a finite element model for mantle dynamics, both in a spherical domain. The continuity of velocities and tractions and the no-slip conditions are imposed on the interfaces. The benchmarks against analytic solutions to the bending of a thin plate verifies that SNAC can give accurate solutions for the traction boundary conditions and that Pyre can correctly handle the crucial data exchanges. As a preliminary model in anticipation of more realistical problems, an elasto-visco-plastic lithosphere is coupled to a Newtonian viscous mantle. Flow in the mantle is driven by a hot

[1]Parts of this work were published by Eh Tan, Eun-seo Choi, Pururav Thoutireddy, Michael Gurnis, and Michael Aivazis in *Geochemistry, Geophysics, Geosystems*, 2006, 7, doi: 10.1029/2005GC001155.

sphere placed at the center of the mantle. This coupled model shows a steady growth of a dome in the lithosphere directly above the hot sphere, as expected. Unnecessarily high numerical cost is caused by the mismatch of time marching schemes adopted in each code.

5.2 Introduction

The dynamic coupling of the crust with the mantle on geological time-scales (such as those occurring on $>10^5$ yrs) is a long-standing problem that has been difficult to approach numerically. Together, the crust and mantle have a complex, composite rheology in which different layers and domains deform on vastly different time and spatial scales (e.g., Goetze and Evans, 1979; Kohlstedt et al., 1995; Ranalli, 1995). The mantle is usually treated as a linear (or non-linear) viscously dominated fluid with an infinite Prandtl number (the ratio of the diffusion of momentum to heat) (Davies, 1999; Schubert et al., 2001). The crust and upper lithosphere is commonly subdivided into several layers of elastic, elasto-plastic, and/or viscoelastic solids (e.g., Bodine et al., 1981; McNutt et al., 1988; Buck, 1991; Burov and Diament, 1995; Albert and Phillips, 2002). Modeling of geodynamics usually involves large displacements, rotation and strains and non-linearities such as brittle fracturing and non-linear viscoelastic creep (e.g., Lavier et al., 2000; Hall et al., 2003). In recent years the adoption of new numerical techniques (e.g., Cundall, 1989; Poliakov et al., 1993a; Fullsack, 1995; Beaumont et al., 2000; Hansen, 2003; O'Neill et al., 2006) has allowed for the modeling of a wide variety of problems related to deformation of lithosphere in 2-D and 3-D. These choices of non-linear rheologies to model the deformation of the mantle and crust as well on the widely different length scales of deformation, typically >10 km for the mantle and from a few meters (faults) to >100 km (long term elastic strength) for the lithosphere poses a vast challenge to numerical modeling of dynamics. The high resolution required for the solutions for 3D multi-scale phenomena leads to exorbitant numerical cost.

Various geophysical problems can be better addressed when the distinct rheolo-

gies of the crust and mantle are modeled as a composite system. For example, the initiation and the evolutionary processes of the Afar triple junction is the result of the mechanical interaction between the lithosphere and a mantle plume (e.g., Manighetti et al., 2001a,b). With a linear viscous rheology only, the transition from continental rifting to the formation of a mid-oceanic ridge through extension of the lithosphere cannot be properly modeled (e.g., Lavier and Manatschal, 2006). On the other hand, without the dynamic motion of the mantle, the time evolution of a triple junction cannot be realistically considered. A more complex example is the Cenozoic evolution of the Basin and Range Province in the western U.S. during which low angle normal faults accommodated extension occurring over a broad region (e.g., Wernicke, 1981). The Basin and Range Province is suggested to be extending under its own weight as well as due to contributions from other mantle-originated forces (e.g., Sonder and Jones, 1999, Chapter 2). The long-term deformation of such regions would be better simulated with a coupling between the mantle and the crust.

In this chapter, we explore the coupling of distinct numerical models as a potential option for such multi-material and multi-scale geodynamic problems. In particular, using facilities provided by Pyre (Tan et al., 2006), one mantle convection code is coupled with a code for modeling lithospheric deformation. We use a variation of CitcomS (**C**alifornia **I**nstitute of **T**echnology **CO**nvection **M**odel, **S**pherical) (Zhong et al., 2000) for mantle convection in a regional spherical cap (RegionalCitcomS hereafter). It is an Eulerian finite element program for solving the Stokes and energy equations and has been widely used for modeling mantle convection. SNAC (**S**tGermai**N** **A**nalysis of **C**ontinua) is used for lithospheric deformation. This code is based on the same algorithm for FLAC (**F**ast **L**agrangian **A**nalysis of **C**ontinua) and the code PARAVOZ developed by Y. Podlatchikov and A. Poliakov (e.g., Poliakov et al., 1993a) but has been extended to 3-D through a computational framework called StGermain (Quenette et al., 2005). SNAC uses an explicit scheme to acquire static or quasi-static equilibrium through a dynamic relaxation technique. Since SNAC can model various rheologies including plasticity, this code is well suited for the modeling of lithospheric deformation, such as those problems studied with the code PARAVOZ or

FLAC (Cundall, 1989; Poliakov et al., 1993a; Burov and Poliakov, 2001; Lavier and Buck, 2002; Hall et al., 2003; Buck et al., 2005). More theoretical background on SNAC is given in Appendix A. In Chapters 3 & 4 we give detailed applications to lithospheric deformation using SNAC without coupling.

5.3 Numerical method

5.3.1 SNAC: Overview

SNAC is a Lagrangian finite difference code for modeling a finitely deforming elasto-visco-plastic solid in 3-D. SNAC is very similar to a conventional finite element code with linear tetrahedral elements and thus shares many benefits with it. However, SNAC does not make explicit use of shape functions. For this reason, it is technically classified as a finite difference code. Another important characteristic of SNAC is that the dynamic relaxation technique is used to achieve static or quasi-static solutions from a dynamic momentum equation. This technique lets SNAC obtain solutions without solving a large matrix equation. A static finite element code whould typically involve inversion of a stiffness matrix and often suffers slow convergence when a non-linear constitutive law is adopted. Since the equation for heat energy balance can be optionally solved, SNAC can solve thermomechanical problems and have temperature-dependent rheologies. Although formulated in the small strain limit, long-term geodynamic problems eventually develop finite deformations. Two potential problems associated with large deformations acquired by accumulating small strains are the loss of objectiveness (the invariance under rotation (e.g. Rudnicki and Rice, 1975)) and the mesh distortion. To avoid these problems, SNAC updates stresses based on the objective Jaumann stress rate (e.g., Rudnicki and Rice, 1975) and performs remeshing when the mesh becomes too distorted.

5.3.2 StGermain framework

SNAC is built up from the StGermain framework, an environment for the development of numerically intense scientific codes (Quenette et al., 2005). As frameworks, StGermain and Pyre differ in that StGermain provides the infrastructural components needed to build complete codes while Pyre provides the superstructure to couple codes together; for this reason we refer to the former as an *infrastructure framework* and the latter as a *superstructure framework*.

StGermain provides a suite of libraries needed by general physical modeling software such as flow control, domain discretization, and initial/boundary condition management (Quenette et al., 2005). One of the advantages of using the StGermain framework is that a code can be easily extended through entry points and plug-ins. By adding or removing entry points, a problem-specific algorithm can be implemented while keeping the modification of source codes limited. Plug-ins are small pieces of codes that work as a part of the program, but can be compiled separately and dynamically loaded at run time. Thus, their development can be completely independent of the main program. For example, a new constitutive relation or time-varying boundary conditions can be made available as plug-ins without modifying or compiling the whole program that already exists.

5.4 Coupling through Pyre

We extend Pyre, a Python-based modeling framework that is originally developed to solve fluid structure interaction problems (Cummings et al., 2002). Technical details as well as an overview of Pyre have been published in (Tan et al., 2006). For completeness, however, we present the most relevant features of this framework and then describe under what physical principles the two solvers, RegionalCitcomS and SNAC, are coupled.

5.4.1 Overview of Pyre

Pyre is a full-featured, object-oriented environment that is capable of handling all aspects of the specification and launch of numerical investigations. Pyre operates on massively parallel supercomputers including both shared memory computers and Beowulf clusters. Pyre is written in the Python programming language, an open-source, well-maintained and widely used interpretive environment.

Pyre leverages the extensibility of the Python interpreter to allow for the seamless integration of rather diverse computational facilities. The framework provides enough flexibility to allow the dynamic discovery of available facilities as part of simulation staging. There is a well-defined and well-documented method by which a new solver or a new material model can be made available to the framework, while the flexibility allows the user to specify solvers and algorithms in the simulation script, without the need for recompilation or relinking. The combination promotes experimentation with new algorithms by lowering the overall overhead associated with trying out new approaches.

Each simulation model under Pyre is called an *Application*. An *Application* could contain one or more *Solvers*. An *Application* and its *Solver(s)* can run on multiple processors, but each processor has only one *Application* and one *Solver* on it (Fig.5.1). The role of the *Application* is to assign each processor a *Solver* and orchestrate the simulation staging of the *Solver(s)*, such as initialization (including memory allocation and variable assignment), time marching, and output (Tan et al., 2006).

5.4.2 *Coupler*

"Coupling" means running two spatially overlapping geophysical applications in such a way that interaction occurs through sharing physical quantities (such as velocity and traction) at interfaces. The sharing of information at an interface typically takes the form of sending and receiving data between two solvers and is managed by a *"Coupler"*.

The *Coupler* is one of the Pyre's classes that control the general procedure of

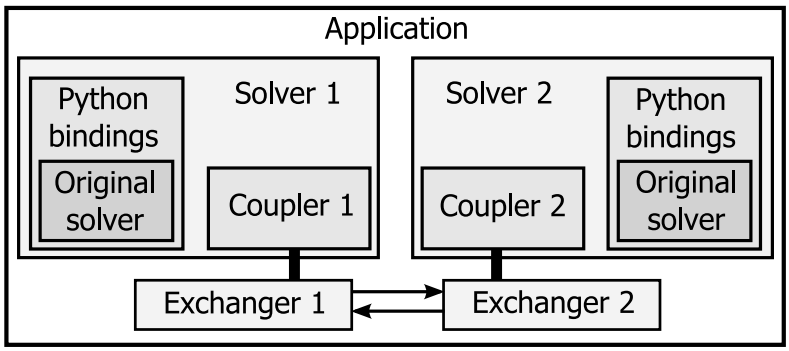

Figure 5.1: The architecture of a coupled *Application*. The *Application* has two *Solvers*. The original code of the solver (in C/C++/Fortran) is complied into a library, which is called by the Pyre *Solver* via the Python bindings. *Solver 1* is the containing solver and has a *Coupler 1*, and the other solver, *Solver 2*, has a *Coupler 2*. The *Couplers* communicate via the *Exchangers*, which are external to the *Solvers*.

coupling between elements in the Pyre. The *Coupler* controls only the top level of coupling, i.e., orchestrates the timing for sending and receiving information in the context of global time marching. The interaction details (for example, what quantities are exchanged and how received information will be used) vary with the nature of the problem (e.g., coupling earthquake rupture to tectonic loading versus coupling crustal deformation with mantle flow).

The actual information exchange occurs in the Exchanger (Fig.5.1), which consists of a number of C++ classes. The *Exchanger* of a *Solver* can communicate with another *Exchanger* of a different *Solver*. An *Exchanger* is specific to its host *Solver*, but independent from the *Solver* that it is coupled to. The different types of the coupling mechanism can be easily implemented by virtue of the flexibility and the extensibility of the *Exchanger*. For example, the *Exchanger* of RegionalCitcomS was originally developed to couple with another CitcomS (Tan et al., 2006); it is used here to couple with SNAC with only little modifications. Since a goal is to leverage existing modeling code, the *Exchanger* is external to the *Solver* and not required for uncoupled applications.

For simplicity, let us first consider the case of a single-processor *Solver* coupled with another single-processor *Solver*. *Solver A*, which is going to send a message, has an *Outlet*, while *Solver B*, which is going to receive, has an *Inlet*. Their meshes

would look like Fig.5.2 or Fig.5.3A in more general settings. First, *Solver B* has a *BoundedMesh*. The *BoundedMesh* contains a set of nodes at the interfaces of coupling *Solvers* and maintains a bounding box of those nodes, hence its name. Here we use "interface" in a loose sense. The set can be the whole collection of boundary nodes of a *Solver* (Fig.5.3B), or only part of it, or the nodes in the overlapping region (Fig.5.3C). The *Inlet* sends the *BoundedMesh* to the *Outlet*. The *Outlet* uses the bounding box as an efficient check on whether the *BoundedMesh* overlaps with the domain of *Solver A*. The *Outlet* then assembles the requested data (usually by interpolation of a local field variable to the nodes in the *BoundedMesh*) and sends them to the *Inlet*. With finite elements, the interpolation involves finding the corresponding element and computes the shape functions of each node in the *BoundedMesh* and is the most time-consuming procedure. If both meshes are static (Eulerian), this procedure is performed only once and the shape functions are stored for subsequent use. If one of the meshes changes with time (i.e., Lagrangian), this procedure repeats at every time step. The *Inlet* then imposes the data received to the interface nodes. Depending on the use of the data, the action of "impose" can have different meanings. If it is used as boundary conditions (BC), the BC arrays are updated; otherwise, it might simply replace the field variable. In the case of multi-processor coupling, the procedure becomes more complicated (for actual codes, see Tan et al., 2006). Each processor still has an *Inlet* (for *Solver B*) or an *Outlet* (for *Solver A*). Additionally, each processor of *Solver A* has a *Source*, but only the leading processor of the *Solver B* has a *Sink*. Each processor of *Solver B* constructs a *BoundedMesh* according to its local mesh. Those local *BoundedMeshes* are broadcast out by the *Sink* to all *Sources*. Each *Source* passes the received *BoundedMesh* coordinates to the *Outlet*, which performs the same interpolation procedure as the single-processor case. The *Outlet* passes the interpolated results to the *Source*. The *Sink* collects the results from all *Sources* in *Solver A* and distributes the collected data to all *Inlets*. *Inlets* then impose these data to the interface nodes. In general, *Solvers* have their own domain decomposition scheme, and the decomposition boundaries of two *Solvers* do not coincide. Therefore, nodes in a local *BoundedMesh* of an *Inlet* might be interpolated by different *Outlets*.

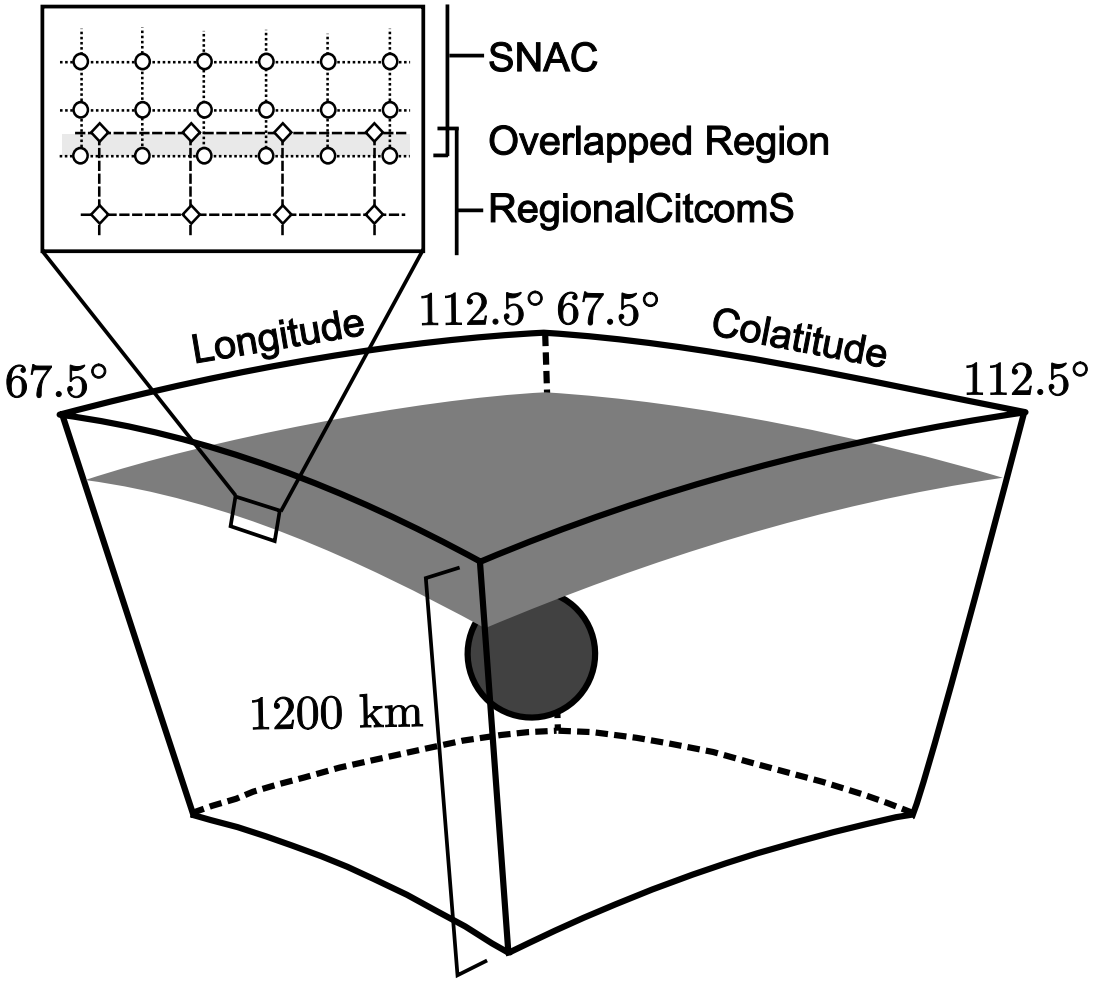

Figure 5.2: Drawings showing the domains of the coupled solvers as a whole. The gray plane at a constant depth represents the interface between two solvers. The spherical object below the interface represents a hot sphere that drives mantle flow.

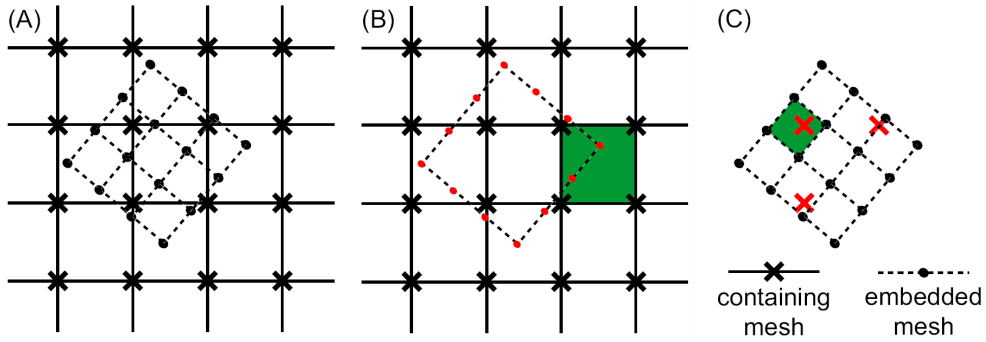

Figure 5.3: An example of a 2-D mesh and a portion of the other mesh. (A) One fine grid is in dashed line, and the other coarse grid in solid line. The fine-grid nodes are shown as dots, and the coarse-gird nodes as crosses. The meshes reside on separated *Solvers* but overlap in the modeling space. The fine mesh is completely immersed within the coarse mesh in this particular example. The two meshes are not required to be parallel to each other. Two scenarios of *BoundedMesh* are presented. (B) The fine-grid *Solver* is *Solver B*. Its *BoundedMesh* consists of twelve boundary nodes (red). The coordinates of these nodes are sent to *Solver A* (the coarse-grid *Solver*) to find the corresponding elements (green, only one element is colored) and shape functions. (C) The coarse-grid *Solver* is *Solver B*. Its *BoundedMesh* is the three nodes (red) in the overlapping region. One corresponding element is shown in green.

The *Sources* and *Sink* maintain the bookkeeping of overlapping nodes.

Different *Solvers*, depending on their design, usually have different coordinate systems and units to represent the physical quantities internally. To facilitate information exchange, we require that any quantities be exchanged in Cartesian coordinates and SI units. Conversion from and to the native coordinate system and units is carefully handled within the *Inlet* and *Outlet*. An option of skipping conversion is available if the *Solvers* use the same coordinate system and units.

During a coupled computation, the *Coupler* monitors the model times of both *Solvers* (for actual codes, see Tan et al., 2006). If the model times of both *Solvers* are equal, they are synchronized. For example, in Fig.5.4, step M+3 and step N+1 are synchronized, but step M+2 and step N+1 are not. Only when the times are synchronized, is the *coarse-grid Solver* allowed to march forward to the next time step. Generally, the *coarse-grid Solver* has a larger stable time step than the *fine-grid Solver*. As a result, at the end of a time step, the *coarse-grid Solver* will be ahead of the *fine-grid Solver* (Fig.5.4A). The *coarse-grid Solver* must wait until the *fine-grid*

Solver catches up (Fig.5.4B-C). The *fine-grid Solver*, if necessary, will clip the size of its stable time step so as to synchronize with the *coarse-grid Solver* (Fig.5.4D).

5.4.3 Coupling RegionalCitcomS and SNAC

The physical principles applied to the interfaces of RegionalCitcomS and SNAC are the continuity of velocity and traction. We also assume the no-slip condition, requiring the continuity of all three components of velocities. A coupling scheme satisfying those principles is implemented as follows:

1. RegionalCitcomS and SNAC are both initiated on a multi-processor computer.
2. RegionalCitcomS gets a solution for an initially given thermal anomaly in the mantle and boundary conditions that are updated later.
3. RegionalCitcomS sends the new stress field to SNAC.
4. SNAC computes a solution (velocity, stress and pressure) for the new traction boundary conditions retrieved from the received stresses.
5. SNAC sends back to RegionalCitcomS an updated velocity field.

Steps 2-5 are repeated until the end of simulation. Since velocities and tractions are determined by one solver at a time and the other simply uses them, the continuity across the interfaces is automatically enforced at all times. This method of coupling a Lagrangian code with an Eulerian code can prevent numerical instabilities caused by substantially different physics working in each code (Fedkiw, 2002).

The time step size for RegionalCitcomS is set to be larger than that of SNAC, which is consistent with the time scales associated with mantle and crustal deformations as well as with implicit and explicit time marching in each code. For this reason, the step (4) includes an inner time loop for SNAC to catch up with RegionalCitcomS. Related to this property, SNAC's dynamic relaxation can potentially pose a problem. A system modeled by SNAC would not be in equilibrium immediately following the update of traction boundary conditions as it takes a finite amount of time to dissipate inertial

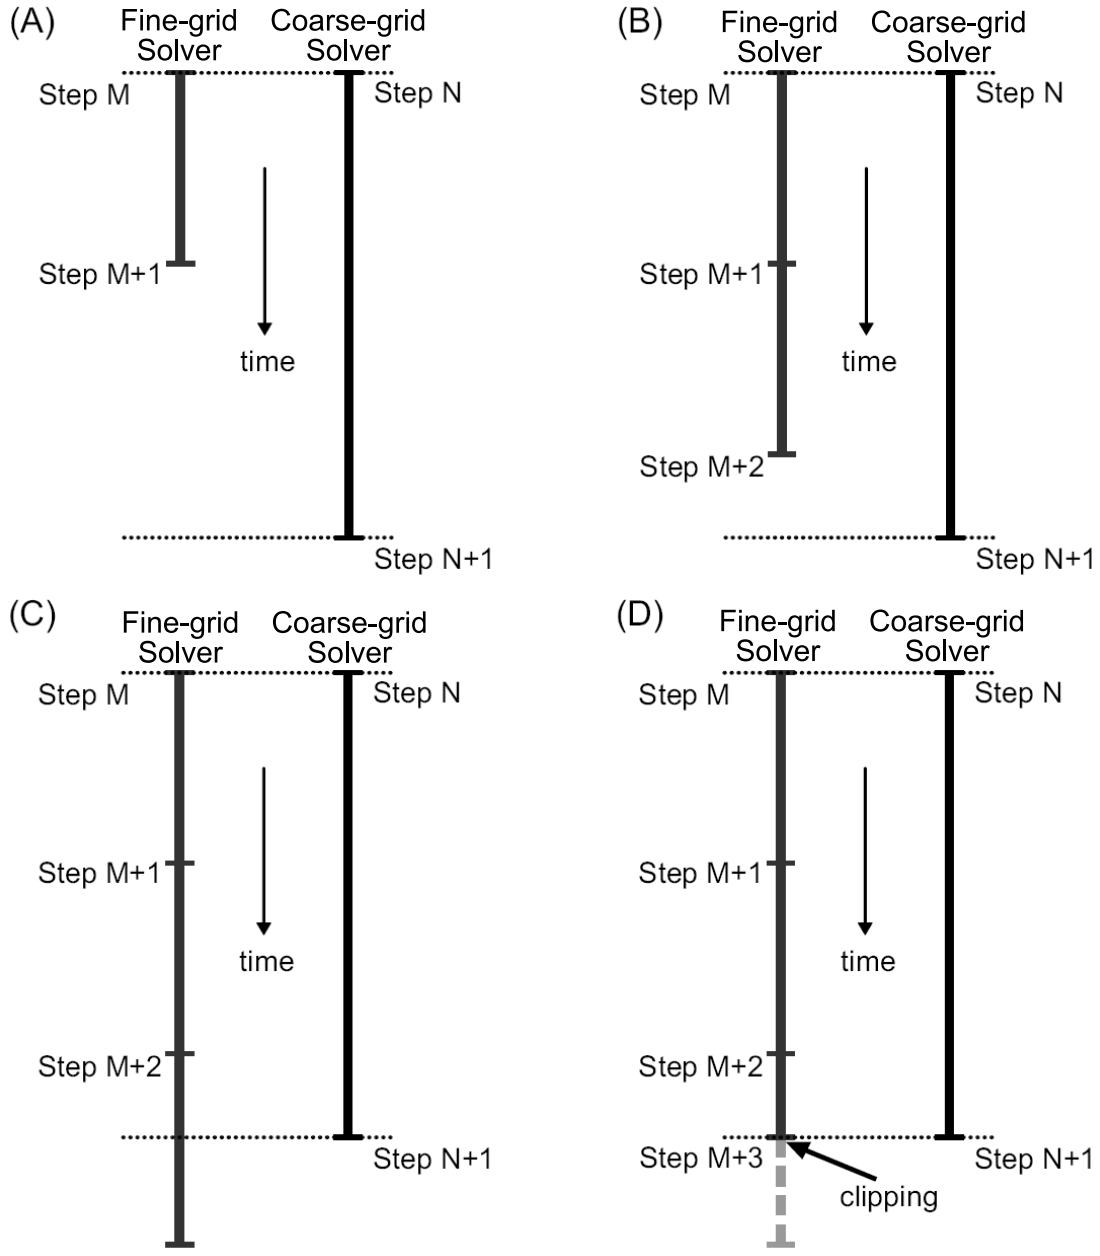

Figure 5.4: Synchronizing time steps of two *Solvers*. (A) After a synchronized step, the coarse-grid *Solver* will be ahead of the fine-grid Solver. (B) and (C) The fine-grid *Solver* keeps marching forward. The coarse-grid *Solver* waits until the fine-grid *Solver* catches up. (D) The fine-grid *Solver* clips the size of its stable time step so as to synchronize with the containing *Solver*.

forces and reach a new equilibrium state within SNAC. Unlike in a true dynamic problem, these transient solutions are not physically meaningful. If a data exchange occurs during this transient phase, RegionalCitcomS would attempt to adjust to that state and can become numerically unstable. To prevent this numerical artifact, transient effects in SNAC should diminish substantially by the time of the next data exchange. In practice, this is achieved by setting the ratio of RegionalCitcomS's time step size to that of SNAC to be sufficiently large.

The incompressibility condition is another problem-dependent detail associated with the coupling of RegionalCitcomS and SNAC. For the convergence of RegionalCitcomS, the incompressibility condition should be enforced for the kinematic boundary conditions. However, the velocity field received from SNAC might have non-zero flux leading to numerical instability. The material near the interface is a Maxwell viscoelastic material in SNAC, according to which viscous responses are guaranteed to be incompressible by the formulation, but even a small dilatational component due to elasticity can violate the RegionalCitcomS's incompressibility condition. Thus, SNAC performs the operation that removes a net flux in its velocity field before being sent to RegionalCitcomS.

Interpolation is a necessary operation in the data exchanges because the meshes of the solvers are usually not perfectly identical (Fig.5.2). Unit conversion is also required because variables are non-dimensional in RegionalCitcomS while SNAC's variables have physical dimensions.

The deformation of the interface is assumed to be negligible on the RegionalCitcomS's side although it is a crucial output from the Lagrangian code, SNAC. If a significant deformation of the interface is expected for the Eulerian code, the use of techniques such as ghost fluids or Arbitrary Lagrangian Eulerian would be required to account for the effect of moving boundaries (Fedkiw, 2002; Arienti et al., 2003). However, the deflection of the interface in this study is on the order of a few km, which is small compared to the typical radial dimension of a RegionalCitcomS's domain (>1000 km). Thus, the effects of moving boundary on the flow field in RegionalCitcomS can be safely ignored.

5.5 Benchmarks

5.5.1 Traction boundary conditions applied on SNAC

A problem of plate bending is solved by SNAC under the thin plate approximation. The purpose is to validate SNAC's solution for traction boundary conditions, which are the type used in the coupled modeling. The analytic solution is easily acquired for the case where two ends are embedded and a uniform transverse loading is applied (Fig.5.5a) (e.g., Turcotte and Schubert, 2001):

$$w(x) = \frac{q}{24D}x^2(L-x)^2, \tag{5.1}$$

where q is the applied load with the unit of N/m^2, L is the length of the plate, and D is the flexural rigidity. The flexural rigidity is defined as $Eh^3/(12(1-\nu^2))$, where E is the Young's modulus, h is the thickness of the plate, ν is the Poisson's ratio. We use different resolutions to test the convergence of numerical solutions to the analytic one. The following values of parameters are used: L=1 km, h=20 m, E= 75 GPa, ν=0.25, and q=10 kPa. While adjusting the resolution, the other dimension of SNAC's 3-D domain is determined such that each hexahedral element becomes a cube. The relative error defined as $\|w-w'\|/\|w\|$ is 1.3 % for the model with 401×7 nodes, where w is the analytically computed flexure, and w' is a numerical solution. Fig.5.5b compares numerical solutions for various resolutions with the analytic one. The expected trend of decreasing error with increased resolution is observed.

5.5.2 The RegionalCitcomS-to-SNAC transfer of traction

A similar numerical experiment is set up to test if Pyre can correctly handle the data exchange between RegionalCitcomS and SNAC. The type of problem is the same with the previous section, but the transverse loading is now computed and given by RegionalCitcomS. Although set up in the spherical coordinate system, thin slices along the equator are used as problem domains. Since the domains are then essentially rectangular slabs, the approximate solutions can be directly compared with the thin

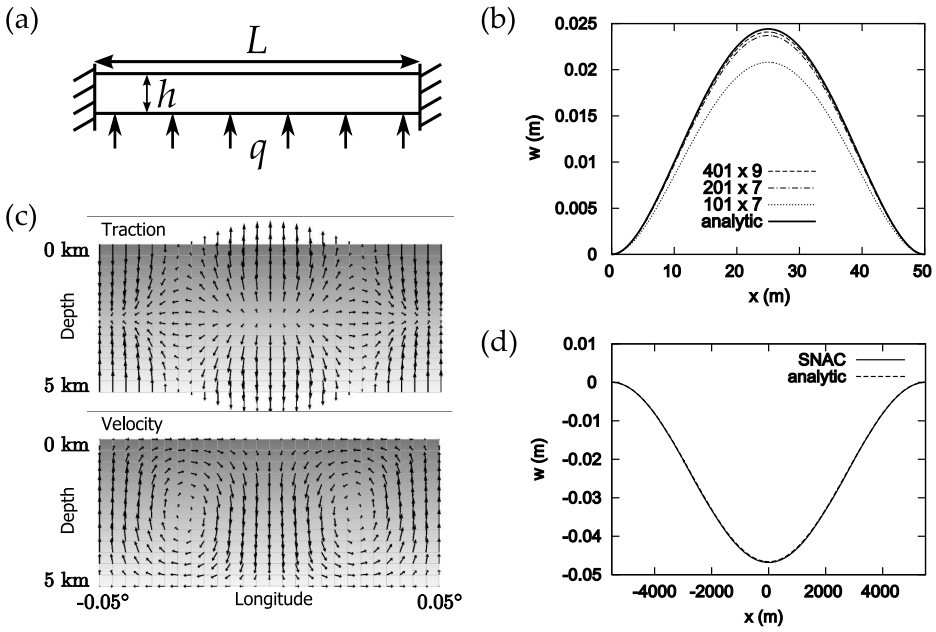

Figure 5.5: (a) A schematic diagram showing the geometry for the thin-plate benchmark problem. (b) Displacement (w) as a function of the transverse distance (x). (c) Arrows represent tractions (upper panel) and velocities (lower panel) resulted from a temperature perturbation with a cosine form. Gray shades represent the temperature normalized by 1700 °C: The darkest for 0 and the lightest for 0.82. (d) Same as (b).

plate theory based on the Cartesian coordinate system. The longitudinal range of the SNAC domain is 0.1° (L=∼10 km) centered on 90° and the thickness is 200 m. The length to thickness ratio is about 50, a sufficiently high value for the thin plate approximation. The RegionalCitcomS's domain covers the same longitudinal range but its radial dimension is 5 km. The overlapped region between the two domains is 15 m.

A cosine perturbation is added to a temperature field of RegionalCitcomS that is otherwise horizontally uniform and varies linearly along the radial direction. Such a temperature field can be described by the following function:

$$T(\varphi, r) = \Delta T_0 \sin^2\left(\pi \frac{r - r_i}{r_o - r_i}\right) \cos\left(2\pi \frac{\varphi - \varphi_{min}}{\varphi_{max} - \varphi_{min}}\right) + (T_t - T_b)\frac{r - r_i}{r_o - r_i}, \qquad (5.2)$$

where r and φ are the radial and longitudinal coordinates, ΔT_0 is the magnitude of perturbation, r_i and r_o are the inner and outer radius of the domain, φ_{min} and

φ_{max} are the minimum and maximum of the longitudinal coordinate of the domain, and T_t and T_b are the top and the bottom temperature. With a uniform viscosity of 10^{21} Pa·s, the resultant velocities and stresses have the form of a cosine function (Fig.5.5c). Then, the traction boundary conditions derived from such a stress field also takes a cosine form. The tractions shown in Fig.5.5c are defined as $\sigma_{ij}(n_r)_j$, where σ_{ij} is the stress tensor and n_r is the unit vector along the radial direction, and computed at each grid point of the RegionalCitcomS domain. Note that the tractions actually applied on the SNAC's boundary are appropriately determined by the local unit normal vector. The Rayleigh number and the values of other parameters are listed in Table 5.1. The velocity boundary conditions for RegionalCitcomS are free-slip on all the surfaces. Zero heat fluxes are assumed on the side walls while the top and the bottom temperature (non-dimensional) are fixed at 0 and 0.82, respectively. The boundary conditions for SNAC are the same with those in the previous section.

A simple analytic solution is acquired for a thin plate with a sinusoidal loading. If the loading is given by $q = q_0 \cos(2\pi x/L)$, the corresponding analytic solution for a thin plate with both ends embedded is

$$w(x) = \left(\frac{L}{2\pi}\right)^4 \frac{q_0}{D}\left[\cos\left(\frac{2\pi}{L}x\right) - 1\right]. \tag{5.3}$$

The numerical solution from a model with 601 by 13 nodes is compared against to the analytic solution in Fig.5.5d. The relative error for this model, defined in the previous section, is 0.62 %.

5.6 Application

5.6.1 Model Setup

As an example of the coupling of RegionalCitcomS and SNAC via Pyre, we solve the following problem motivated by the thermal and mechanical response of continental lithosphere to a persistent hot spot. A domain for viscous mantle is defined by lines

Table 5.1: List of parameters

For the benchmarks	
Ra	4.226363×10^{8}
ΔT	1700 °C
η_{ref}	10^{21} Pa·s
κ	10^{-6} m^2/s
g	9.8 m/s^2
R	6371 km
α	3.0×10^{-5} K^{-1}
ρ	3280 kg/m^3
ΔT_0	10^{-8*}
r_i	0.99921519384711976*
r_o	0.99997017736619054*
φ_{min}	1.5699236621688994
φ_{max}	1.5716689914208937
T_t	0*
T_b	0.82*
For the coupled application†	
ΔT_0	0.18*
d_{blob}	0.07*
E	45.0
C	1.0

*Non-dimensionalized with respect to R or ΔT.

†The Rayleigh number and its defining parameters are the same with those for the benchmarks.

of constant colatitude and longitude, marking 67.5° and 112.5° respectively (Fig.5.2). The top and bottom boundary of the domain are at the depth of 145 km and 1200 km, respectively. A portion of the lithosphere resides in a separate but overlapping domain defined by the lines of the same colatitude and longitude. This lithosphere has a thickness of 150 km, creating a 5 km-thick region overlapped with the mantle (Fig.5.2). Since the overlapped region is sufficiently thin that one can ignore radial gradients of all the variables within the region, it is essentially same as having a single surface as an interface between two solvers. This model was assigned 432 processors, 400 for SNAC and 32 for RegionalCitcomS. The computation took 10 days for the model time of 1.5 Myrs.

The entire system is composed of 100 Ma elasto-visco-plastic lithosphere, Newtonian viscous mantle, and a hot sphere in the mantle. Temperatures are fixed at 1400 °C on the bottom surface of RegionalCitcomS, and the top surface has a fixed temperature corresponding to the depth of 145 km according to a half space cooling model. SNAC has the temperature of 0 °C on the top surface, but the bottom temperature is determined in the same way with RegionalCitcomS (Fig.5.6a). The hot sphere, placed at the depth of 700 km, is defined by the following temperature field:

$$T(\vartheta, \varphi, r) = \Delta T_0 \exp(-(d/d_{blob})^2) + T_{amb}(\vartheta, \varphi, r), \tag{5.4}$$

where ΔT_0 is the temperature difference between the center of the blob and the reference mantle temperature, d is the Euclidean distance from the center of the blob, d_{blob} is the characteristic length scale equivalent to the radius of the blob, and T_{amb} is the ambient temperature field given by a half-space cooling model. Zero heat fluxes are assumed on the side walls of both codes. The radial profiles of temperature through normal mantle and the hot sphere are shown in Fig.5.6c. Table 5.1 lists the values used for these parameters.

Kinematic boundary conditions are free-slip on all the surfaces of both codes except for the top surface of RegionalCitcomS and the side walls of SNAC. The top surface of the mantle domain has zero velocities in the beginning, but SNAC-derived

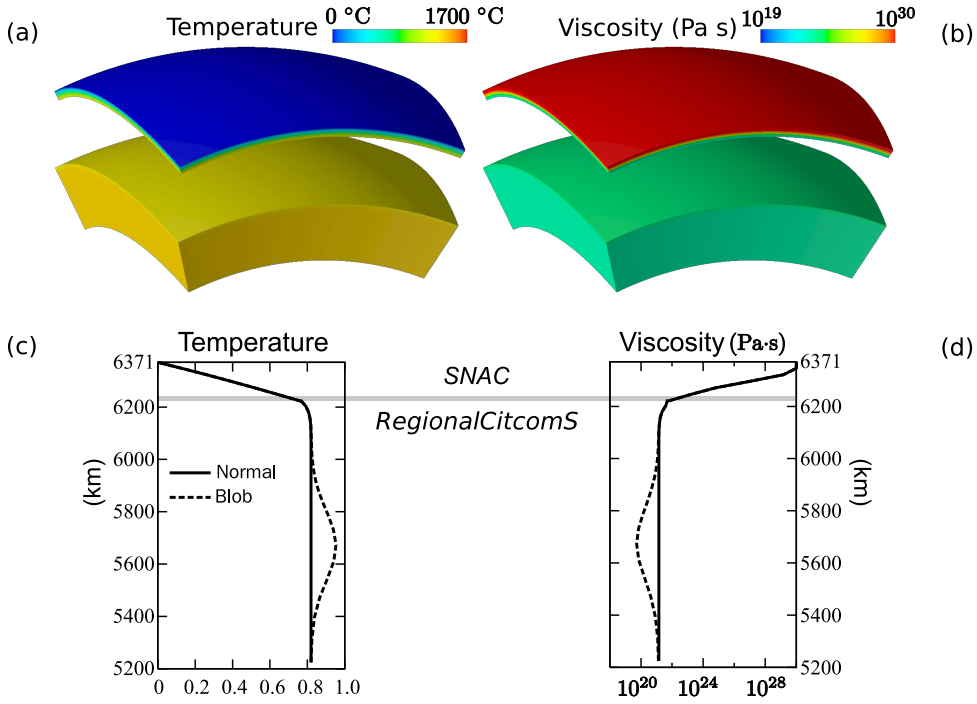

Figure 5.6: (a,b) Temperature and viscosity distribution mapped on the 3-D domains of each solver. (c,d) The depth profiles of temperature and viscosity taken right through the hot sphere are compared with the reference profiles. Note that the temperature is normalized with respect to 1700 °C in (c).

velocities are imposed on it later. The SNAC's side walls are all fixed like the embedded plate discussed in the previous section. As described earlier, SNAC has traction boundary conditions for the bottom surface (the interface with the mantle solver), derived from the stress field received from RegionalCitcomS.

Viscosity is given by a temperature-dependent function:

$$\eta(T) = \eta_{ref} \exp\left(\frac{E}{T+C} - \frac{E}{1.0+C}\right), \tag{5.5}$$

where η_{ref} is the reference viscosity, T is the non-dimensional temperature, E and C are constants determining the range of viscosity variations. In this study parameters are chosen such that the viscosity varies from 10^{21} to 10^{28} Pa·s. This function is shared by both codes so that the common temperature distribution can ensure the continuity of viscosity across the interface (Fig.5.6b,d). The rheology for lithosphere is elasto-visco-plastic with the Mohr-Coulomb yield criterion and a non-associated flow rule (Albert et al., 2000; Poliakov et al., 1993b; Simo and Hughes, 2004). All the parameters and their values are listed in Table 5.1.

5.6.1.1 A qualitative 3-D test

This test verifies that the 3-D coupled model behaves reasonably in a qualitative sense. For simplicity, SNAC represents an elastic plate while mantle flow in RegionalCitcomS is driven by a hot blob as described in the previous section. However, the sphere's depth and radius are slightly different such that it is located at the depth of 573.39 km and the radius is about 400 km. The elastic flexure computed by the coupled SNAC is compared with topographies computed by the standalone RegionalCitcomS. We obtain topography from RegionalCitcomS in two different ways: One is to convert the σ_{rr} component of stress tensor to topography under the assumption of Airy isostasy; the other is to directly integrate over time the radial component of velocity on the top surface (Zhong et al., 1996). Since pressure from the evolving topography is applied on the top surface as a traction boundary condition, the system eventually reaches an equilibrium state in which the internal loading is perfectly balanced with the load of

topography. Since the growth of topography is much faster than the advection of the hot sphere, we can obtain a surface topography that is almost time independent and comparable to the one by stress conversion. In general, the two types of calculation should coincide (Zhong et al., 1996).

We compare various profiles of the surface topography along the equator from the coupled SNAC and the standalone RegionalCitcomS in Fig.5.7. Since we are interested only in the instantaneous elastic response, it is sufficient to let SNAC and RegionalCitcomS interact only once in the beginning. The inertial forces in SNAC were sufficiently damped out after 2000 time steps. So, the topography from the coupled SNAC is constant afterward and it is this static topography that we use for the comparison. Consistent with physical intuition, the elastic plate added on the top of RegionalCitcomS reduces the height of the central dome when compared with that achieved by RegionalCitcomS alone. Furthermore, the height is inversely proportional to flexural rigidity of the plate.

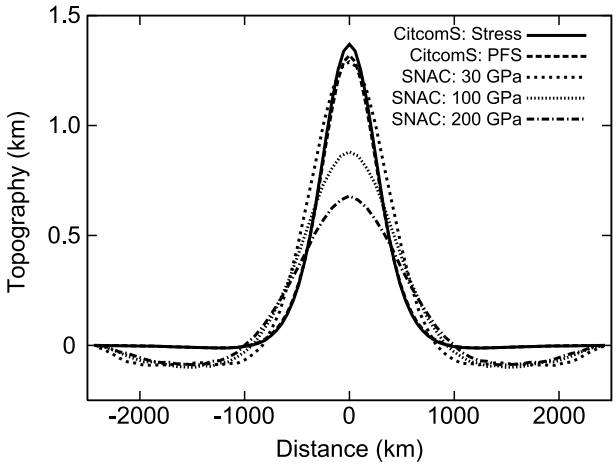

Figure 5.7: Topography along the equator for various coupled models.

5.6.2 Results

The overall behavior of the coupled model can be described by a dome growing in response to a hot sphere in the mantle. The temporal evolution of topography is tracked at the top center of the SNAC's domain (Fig.5.8a). The lithosphere with

no pre-stress responded to the mantle flow such that the topography grew to about 380 m during the first 10^4 yrs. Afterward, the viscoelastic plate is slowly flexured further upward, reaching a topography of about 520 m at 1.5 Myrs. Magnitudes of the radial tractions that drive the deformation of lithosphere are in the order of 1 GPa (Fig.5.8b). The magnitudes increase with time as the rising hot sphere gets near the interface. The rising speed of the hot sphere is about 10 cm/yr but the magnitude of velocity near the interface is must smaller ($\leq\sim$1 mm/yr) because the lithosphere is elastically stiff, more viscous, and initially at rest (Fig.5.8c). After adjusting to the velocities received from SNAC for the first time at 0.2 Myr, RegionalCitcomS reaches a steady state. The same trends in topography, traction, and velocity are depicted in Fig.5.9. This figure, however, shows profiles of each variable along the equator from 0 to 1.5 Myr at the 0.3 Myr interval.

5.6.3 Discussion

According to the above results, the coupled model could simulate a continuous and predictable process of doming due to a thermal anomaly in the mantle. Although the rheology is capable of non-linear processes such as yielding and localization of plastic strain, the maximum amount of deformation was not sufficient for the triggering of these processes. The computation took 10 days for the model time of 1.5 Myrs using 432 processors. With the same problem sizes and the number of processors, a simple extrapolation gives an estimate that two months would take to cover about 10 Myrs of model time, which is a prohibitively long time considering large parameter spaces to explore.

The rate-limiting factor in this coupling study is the high ratio between two solver's time step size. Currently, that ratio of SNAC to RegionalCitcomS is 10^4. Such a high ratio was required because of the dynamic nature of SNAC as discussed in 5.4.3. If SNAC were an implicit code, the whole computation would be greatly accelerated. Although each time step would be more expensive than that of the current explicit time marching, the convergence to a certain static or steady state would be generally faster

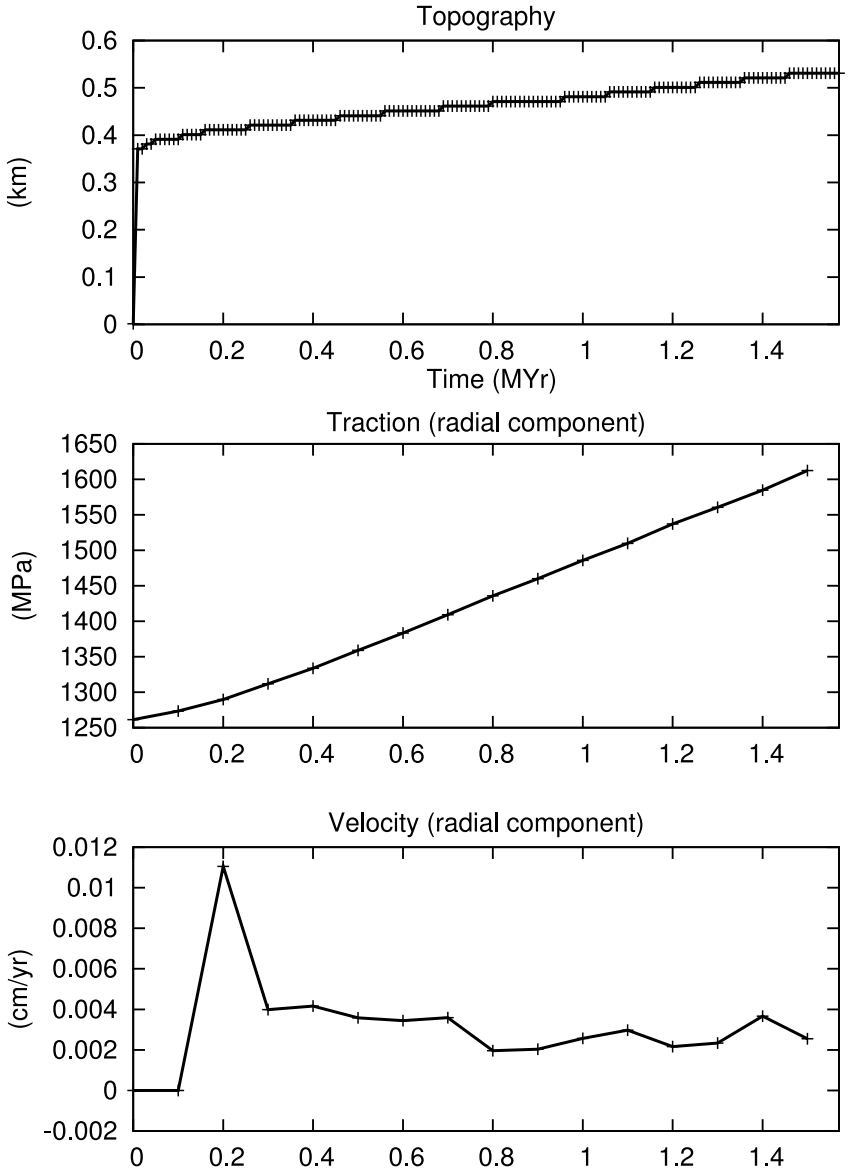

Figure 5.8: Temporal variations of topography from SNAC (top), the radial component of traction from RegionalCitcomS (middle), and the radial component of velocity of RegionalCitcomS (bottom). All three quantities were taken at the top center of each solver's domain.

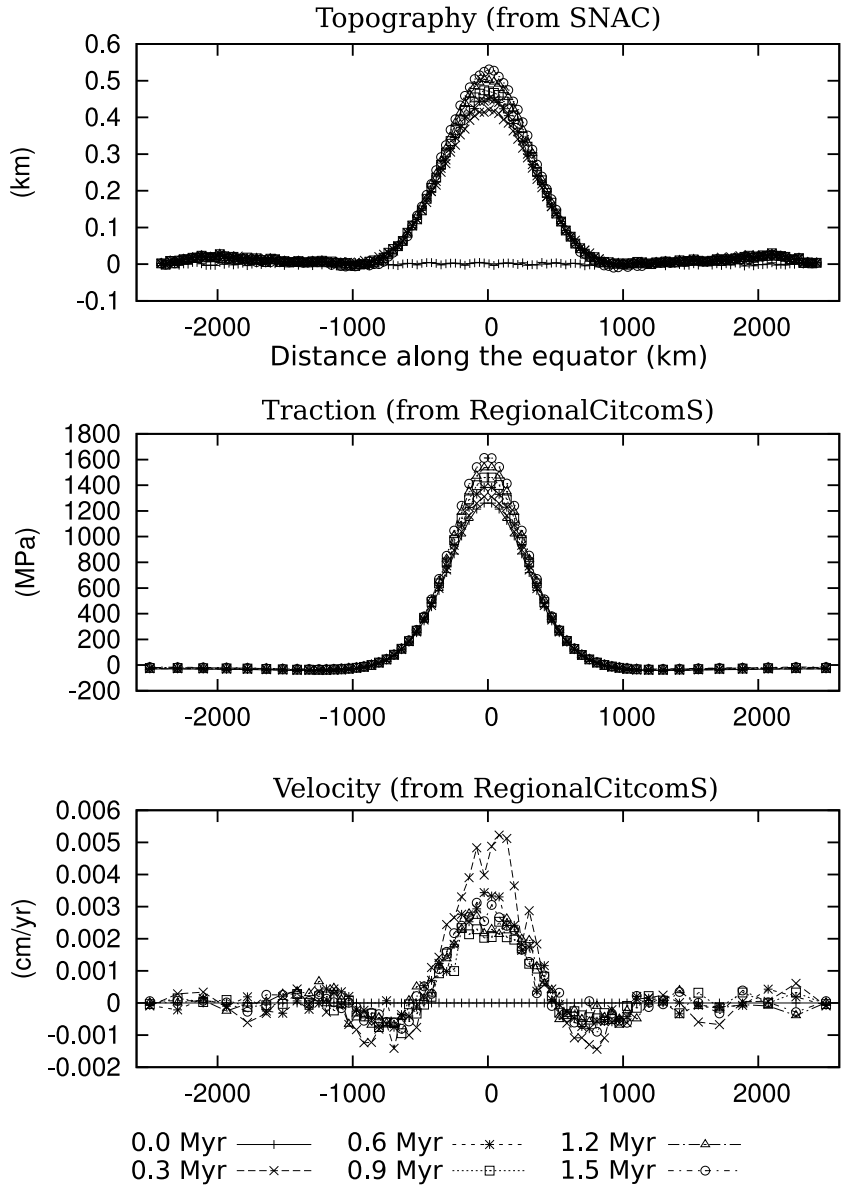

Figure 5.9: Along-equator profiles of topography and the radial component of traction and velocity.

than damping artificial inertial forces. The efficiency in parallelism implemented in Pyre as well as within each node can be another issue, but it is only secondary in the present case. If this demanding numerical cost can be reduced, the coupling technique presented in this chapter can be applied to a variety of geophysical problems.

References

Albert, R., Phillips, R., Dombard, A., Brown, C., 2000. A test of the validity of yield strength envelope with an elastoviscoplastic finite element model. Geophys. J. Int. 140, 399–409.

Albert, R. A., Phillips, R. J., 2002. Time-dependent effects in elastoviscoplastic models of loaded lithosphere. Geophys. J. Int. 151, 612–621.

Arienti, M., Hung, P., Morano, E., Shepherd, J. E., 2003. A level set approach to Eulerian-Lagrangian coupling. J. Comput. Phys. 185, 213–251.

Beaumont, C., Munoz, J. A., Hamilton, J., Fullsack, P., 2000. Factors controlling the Alpine evolution of the central Pyrenees inferred from a comparison of observations and geodynamical models. J. Geophys. Res. 105, 8121–8145.

Bodine, J. H., Steckler, M. S., Watts, A. B., 1981. Observations of flexure and the rheology of the oceanic lithosphere. J. Geophys. Res. 86, 3695–3707.

Buck, W., Lavier, L., Poliakov, A., 2005. Modes of faulting at mid-ocean ridges. Nature 434, 719–723.

Buck, W. R., 1991. Modes of continental lithospheric extension. J. Geophys. Res. 96, 20,161–20,178.

Burov, E., Poliakov, A., 2001. Erosion and rheology controls on synrift and postrift evolution; verifying old and new ideas using a fully coupled numerical model. J. Geophys. Res. 106, 16,461–16,481.

Burov, E. B., Diament, M., 1995. The effective elastic thickness (T_e) of continental lithosphere; what does it really mean? J. Geophys. Res. 100, 3905–3927.

Cummings, J., Aivazis, M., Samtaney, R., Radovitzky, R., Mauch, S., Meiron, D., 2002. A virtual test facility for the simulation of dynamic response in materials. J. Supercomput. 23, 39–50.

Cundall, P., 1989. Numerical experiments on localization in frictional materials. Ingenieur Archiv. 58, 148–159.

Davies, G. F., 1999. Dynamic Earth: Plates, Plumes, and Mantle Convection. Cambridge University Press, Cambridge.

Fedkiw, R., 2002. Coupling an Eulerian fluid calculation to a Lagrangian solid calculation with the Ghost Fluid Method. J. Comput. Phys. 175, 200–224.

Fullsack, P., 1995. An arbitrary Lagrangian-Eulerian formulation for creeping flows and its application in tectonic models. Geophys. J. Int. 120, 1–23.

Goetze, C., Evans, B., 1979. Stress and temperature in the bending lithosphere as constrained by experimental rock mechanics. Geophys. J. Roy. Astron. Soc. 59, 463–478.

Hall, C., Gurnis, M., Sdrolias, M., Lavier, L. L., Mueler, R. D., 2003. Catastrophic initiation of subduction following forced convergence across fracture zones. Earth Planet. Sci. Lett. 212, 15–30.

Hansen, D. L., 2003. A meshless formulation for geodynamic modeling. J. Geophys. Res. 108, 2549, doi:10.1029/2003JB002460.

Kohlstedt, D., Evans, B., Mackwell, S., 1995. Strength of the lithosphere: Constraints imposed by laboratory experiments. J. Geophys. Res. 100 (B9), 17587–17602.

Lavier, L. L., Buck, W., 2002. Half graben versus large-offset low-angle normal fault: Importance of keeping cool during normal faulting. J. Geophys. Res. 107 (B6, 2122), 10.1029/2001JB000513.

Lavier, L. L., Buck, W. R., Poliakov, A. N. B., 2000. Factors controlling normal fault offset in an ideal brittle layer. J. Geophys. Res. 105 (B10), 23,431–23,442.

Lavier, L. L., Manatschal, G., 2006. A mechanism to thin the continental lithosphere at magma-poor margins. Nature 440, 324–328.

Manighetti, I., King, G. C. P., Gaudemer, Y., Scholz, C. H., Doubre, C., 2001a. Slip accumulation and lateral propagation of active normal faults in Afar. J. Geophys. Res. 106, 13,667–13,695.

Manighetti, I., Tapponnier, P., Courtillot, V., Gallet, Y., Jacques, E., Gillot, P. Y., 2001b. Strain transfer between disconnected, propagating rifts in Afar. J. Geophys. Res. 106, 13,613–13,665.

McNutt, M. K., Diament, M., Kogan, M. G., 1988. Variations of elastic plate thickness at continental thrust belts. J. Geophys. Res. 93, 8825–8838.

O'Neill, C., Moresi, L., Müller, D., Albert, R., Dufour, F., 2006. Ellipsis 3d: A particle-in-cell finite-element hybrid code for mantle convection and lithospheric deformation. Comput. Geosci. 32, 1769–1779.

Poliakov, A., Podladchikov, Y., Talbot, C., 1993a. Initiation of salt diapirs with frictional overburdens: Numerical experiments. Tectonophysics 228 (3-4), 199–210.

Poliakov, A. N. B., Cundall, P. A., Podladchikov, Y. Y., Lyakhovsky, V. A., 1993b. An explicit inertial method for the simulation of viscoelastic flow: An evaluation of elastic effects on diapiric flow in two- and three-layers models. In: Flow and Creep in the Solar Systems: Observations, Modeling and Theory. Kluwer Academic Publishers, pp. 175–195.

Quenette, S., Appelbe, B., Gurnis, M., Hodkinson, L., Moresi, L., Sunter, P., 2005. An investigation into design for performace and code maintainability in high performance computing. ANZIAM J. 46 (E), C101–C116.

Ranalli, G., 1995. Rheology of the Earth. Chapman & Hall, London, United Kingdom.

Rudnicki, J., Rice, J., 1975. Conditions for the localization of deformation in pressure-sensitive dilatant materials. J. Mech. Phys. Solids. 23, 371–394.

Schubert, G., Turcotte, D. L., Olson, P., 2001. Mantle Convection in the Earth and Planets. Cambridge University Press, Cambridge.

Simo, J., Hughes, T., 2004. Computational Inelasticity. Springer, New York.

Sonder, L. J., Jones, C. H., 1999. Western United States extension; how the west was widened. Annu. Rev. Earth Planet. Sci. 27, 417–462.

Tan, E., Choi, E., Thoutireddy, P., Gurnis, M., Aivazis, M., 2006. Geoframework: Coupling multiple models of mantle convection within a computational framework. Geochem. Geophys. Geosyst. 7, Q06001, doi:10.1029/2005GC001155.

Turcotte, D. L., Schubert, G., 2001. Geodynamics, 2nd Edition. Cambridge University Press.

Wernicke, B., 1981. Low-angle normal faults in the Basin and Range Province; nappe tectonics in an extending orogen. Nature 291 (5817), 645–648.

Zhong, S., Gurnis, M., Moresi, L., 1996. Free-surface formulation of mantle convection - I. basic theory and application to plumes. Geophys. J. Int. 127 (3), 708–718.

Zhong, S., Zuber, M. T., Moresi, L., Gurnis, M., 2000. The role of temperature-dependent viscosity and surface plates in spherical shell models of mantle convection. J. Geophys. Res. 105 (B5), 11,063–11,082.